

Exact NMR simulation of protein-size spin systems using tensor train formalism

D. V. Savostyanov[†], S. V. Dolgov[§], J. M. Werner[‡], Ilya Kuprov[†]

February 18, 2014

Abstract

We introduce a new method, based on alternating optimization, for compact representation of spin Hamiltonian, and solution of linear systems in the tensor train format. We demonstrate its utility by simulating, without significant approximations, a ^{15}N NMR spectrum of ubiquitin — protein containing several hundred interacting nuclear spins. Existing simulation algorithms for the spin system and the NMR experiment in question either require significant approximations or scale exponentially with the system size. We compare the proposed method to the *Spinach* package that uses heuristic restricted state space (RSS) techniques to achieve polynomial complexity scaling. When the spin system topology is close to a linear chain (e.g. for backbone of a protein), the tensor train representation of a Hamiltonian is more compact and can be computed faster than the sparse representation using the RSS.

Keywords: density matrix renormalization group, alternating minimal energy, tensor train, nuclear magnetic resonance, protein

1 Introduction

The amount of patience required to simulate exactly a nuclear magnetic resonance (NMR) spectrum of an N -spin system scales approximately as $\mathcal{O}(2^N)$. That much is rarely available, and considerable thought has consequently been given over the last decade to more efficient methods of computing NMR spectra [1, 7, 14, 22, 51], in particular to algorithms that promise to achieve that objective in polynomial time. Such algorithms do exist [7, 24], but they make significant a priori assumptions about the spin system evolution — it is usually assumed that the system stays weakly correlated for the duration of the experiment [24, 35].

Outside the NMR community, significant progress was recently made in the development of tensor structured methods [46, 32, 47, 29, 20], all of which descend broadly from the density matrix renormalization group (DMRG) formalism [55,

[†]University of Southampton, School of Chemistry, Highfield Campus, Southampton SO17 1BJ, United Kingdom (d.savostyanov@soton.ac.uk)

[‡]University of Southampton, Centre for Biological Systems, Highfield Campus, Southampton SO17 1BJ, United Kingdom

[§]Max-Planck Institute for Mathematics in the Sciences, Inselstrasse 22, Leipzig 04103, Germany

56] and matrix product states (MPS) [18, 31]. Typical applications of DMRG algorithms are 1D chains at absolute zero [46] and finite temperature [6, 48, 54], with extension to 2D lattices [36, 38, 58, 49]. Invariably, however, the interesting spin systems encountered in the daily practice of biological NMR spectroscopy (proteins, polynucleotides, polysaccharides) are all irregular three-dimensional room-temperature agglomerations with multiple interlocking loops in the spin coupling graph and no identical couplings [9]. When the strict NMR spectroscopy requirement for correct wavefunction phase during the very long (milliseconds to seconds) dissipative NMR spin system trajectories is added to the list, time-domain DMRG methods are currently struggling.

There are some biologically relevant cases, however, that may still be treated as linear chains — for the purposes of simulating simple backbone NMR experiments, protein side chains may often be ignored. This makes the corresponding spin system a weakly branched linear chain that is amenable to DMRG type treatment. Simple NMR experiments can also be reformulated as a matrix-inverse-times-vector problem in the frequency domain, for which efficient algorithms in tensor product formats have recently emerged [23, 11, 12, 13]. We report in this communication the behavior of the AMEn algorithm [12, 13], applied to the solution of the NMR simulation problem in the frequency domain, as well as to the technical task of adding together, without loss of accuracy, tensor train representations of thousands of spin Hamiltonian terms for a protein.

Having integrated the algorithms described below into *Spinach* (a large-scale magnetic resonance simulation library [22]), we are reporting here the first exact quantum mechanical simulation of a liquid-state 1D NMR spectrum for a protein backbone spin system with several hundred coupled spins. Beyond the physical assumptions made by chemists at the problem formulation stage and the controllable numerical rounding error of the tensor train format itself [40], there are no approximations.

2 Tensor product formats

Tensor product expressions appear naturally in spin dynamics because the state space of a multi-spin system is a direct product of the state spaces of individual spins [17]. A simple example is the nuclear Zeeman interaction Hamiltonian

$$\hat{H}_Z = \sum_{n=1}^N \vec{B}_0 \cdot \mathbf{A}^{(n)} \cdot \hat{\mathbf{S}}^{(n)}, \quad (1)$$

where N is the number of spins, \vec{B}_0 is the applied magnetic field, $\mathbf{A}^{(n)}$ are nuclear chemical shielding tensors, and the sum runs over all nuclei. Components of the nuclear spin operators $\hat{\mathbf{S}}^{(n)} = \begin{bmatrix} \hat{S}_X^{(n)} & \hat{S}_Y^{(n)} & \hat{S}_Z^{(n)} \end{bmatrix}$ have the following tensor product form

$$\hat{S}_{\{X,Y,Z\}}^{(n)} = \mathbf{1} \otimes \cdots \otimes \mathbf{1} \otimes \hat{\sigma}_{\{X,Y,Z\}} \otimes \mathbf{1} \otimes \cdots \otimes \mathbf{1}, \quad (2)$$

where $\mathbf{1}$ denotes an identity matrix of appropriate dimension and the Pauli matrices $\hat{\sigma}_X, \hat{\sigma}_Y, \hat{\sigma}_Z$ occur in the n -th position of the tensor product sequence. This representation is known in numerical linear algebra as the canonical polyadic (CP)

format [32]. Although CP representations have been known in magnetic resonance spectroscopy for a long time [5], they suffer in practice from rapid inflation — spin Hamiltonians encountered in NMR and ESR (electron spin resonance) systems can be complicated [17] and, even for simple initial conditions, the number of terms in the canonical decomposition increases rapidly during system evolution. More ominously, the number of CP terms can change dramatically after small perturbations in the Hamiltonian or the system state. A simple example is

$$\hat{S}_Z = \sum_{n=1}^N \hat{S}_Z^{(n)} = \lim_{\varepsilon \rightarrow 0} \frac{(\mathbf{1} + \varepsilon \hat{\sigma}_Z)^{\otimes N} - \mathbf{1}^{\otimes N}}{\varepsilon}, \quad (3)$$

where $\hat{a}^{\otimes N} = \hat{a} \otimes \dots \otimes \hat{a}$. The left hand side of this equation contains N terms, but the expression approximating it on the right hand side has only two terms, and one could be tempted to use it to reduce storage and CPU time. However, both terms of the approximation grow to infinity when $\varepsilon \rightarrow 0$, and the accuracy is lost due to rounding errors. Such instabilities in the CP format make it difficult to use — in finite precision arithmetic the number of terms in the decomposition quickly becomes equal to the dimension of the full state space and any efficiency savings disappear.

Unlike the CP format, which is an open tensor network, *closed* tensor network formats are stable to small perturbations. The most popular closed tensor network format was repeatedly rediscovered and is currently known under three different names: DMRG in condensed-matter physics [55, 56], MPS in computational physics [18, 31, 42], and TT (tensor train) in numerical linear algebra [40]. It is defined as follows

$$\hat{\chi} = \tau(\hat{\chi}^{(1)}, \dots, \hat{\chi}^{(N)}) = \sum_{\alpha_1, \dots, \alpha_{N-1}} \hat{\chi}_{\alpha_1}^{(1)} \otimes \hat{\chi}_{\alpha_1, \alpha_2}^{(2)} \otimes \dots \otimes \hat{\chi}_{\alpha_{N-1}}^{(N)}. \quad (4)$$

The TT representation of the total \hat{S}_Z operator in Eq. (3) is similar to the high-dimensional Laplacian [28]:

$$\hat{S}_Z = \sum_{\alpha_1=1}^2 \dots \sum_{\alpha_{N-1}=1}^2 \hat{H}_{\alpha_1}^{(1)} \otimes \hat{H}_{\alpha_1, \alpha_2}^{(2)} \otimes \dots \otimes \hat{H}_{\alpha_{N-1}}^{(N)}, \quad (5)$$

$$\hat{H}^{(1)} = \begin{bmatrix} \hat{\sigma}_Z & \mathbf{1} \end{bmatrix}, \quad \hat{H}^{(2)} = \dots = \hat{H}^{(N-1)} = \begin{bmatrix} \mathbf{1} & 0 \\ \hat{\sigma}_Z & \mathbf{1} \end{bmatrix}, \quad \hat{H}^{(N)} = \begin{bmatrix} \mathbf{1} \\ \hat{\sigma}_Z \end{bmatrix}.$$

The number of terms in each summation (known as *bond dimension*, or *TT rank*) is two, and all entries of the decomposition are now bounded. The TT representation in the right hand side of Eq. (5) has $4N - 4$ operators, each of which is either zero, or the identity $\mathbf{1}$, or the Pauli matrix $\hat{\sigma}_Z$. The CP representation of \hat{S}_Z in Eq. (3) has N^2 such operators — the tensor train representation is clearly more memory-efficient.

Another notable example is the ZZ coupling Hamiltonian that often makes an appearance in simple linear spin chain models:

$$\hat{J} = \sum_{m>n} \hat{S}_Z^{(n)} \hat{S}_Z^{(m)}. \quad (6)$$

This is a CP expression with $N(N - 1)$ terms and $N^2(N - 1)$ operators. The corresponding TT representation is

$$\hat{J} = \sum_{\alpha_1=1}^3 \cdots \sum_{\alpha_{N-1}=1}^3 \hat{J}_{\alpha_1}^{(1)} \otimes \hat{J}_{\alpha_1, \alpha_2}^{(2)} \otimes \cdots \otimes \hat{J}_{\alpha_{N-1}}^{(N)},$$

$$\hat{J}^{(1)} = \begin{bmatrix} 0 & \hat{\sigma}_Z & \mathbf{1} \end{bmatrix}, \quad \hat{J}^{(2)} = \cdots = \hat{J}^{(N-1)} = \begin{bmatrix} \mathbf{1} & 0 & 0 \\ \hat{\sigma}_Z & \mathbf{1} & 0 \\ 0 & \hat{\sigma}_Z & \mathbf{1} \end{bmatrix}, \quad \hat{J}^{(N)} = \begin{bmatrix} \mathbf{1} \\ \hat{\sigma}_Z \\ 0 \end{bmatrix}. \quad (7)$$

Here each summation runs over three terms only, and the total number of matrices $\hat{J}_{\alpha_{n-1}, \alpha_n}^{(n)}$ is $9N - 12$: much fewer than the storage requirements of the CP format.

The storage requirements of tensor structured representations (both CP and TT) stand in sharp contrast with the classical approach to magnetic resonance simulations [1, 51], where the Hamiltonian is represented as a $2^N \times 2^N$ sparse matrix with all non-zero entries stored in memory. As soon as the sparse matrix is assembled, the CPU and memory resources grow exponentially with the number of spins N , making the simulation prohibitively difficult for large systems. To avoid this problem (known colloquially as the *curse of dimensionality*), all data ideally should be kept in compressed tensor formats of the form given in Eqs. (1) and (4), without ever expanding the Kronecker products. TT format achieves this objective [32, 29, 20].

A very considerable body of literature exists on manipulating expressions directly in tensor product formats. In particular, a given tensor may be converted into the TT format using sequential singular value decompositions [46, 42, 40]. Given tensors in the TT format, one can perform many linear or bilinear operations (addition, element-wise multiplication, matrix-vector multiplication) [47, 40], Fourier transform [10], and convolution [27], avoiding exponentially large arrays and computational costs.

These developments would have permitted large-scale magnetic resonance simulations entirely in the TT format, were it not for a significant obstacle — the summation operation in tensor train representation is an expensive procedure that carries a significant accuracy penalty due to a large number of tensor truncation (rounding) steps. This was also previously noted for three-dimensional potentials occurring in electron density computations [30, 45, 19]. Spin Hamiltonians of practically interesting biological systems contain many thousands one- and two-spin terms of the kind shown in Eq. (2). Intermediate expressions in spin dynamics simulations also frequently involve large sums. We demonstrate below that in those circumstances the standard bundle-and-recompress tensor network summation procedure leads either to the bond dimension expansion beyond the limits of modern computing hardware, or to a catastrophic accuracy loss. Here we propose an alternative algorithm for computing large sums, based on alternating tensor train optimization, and use it to carry out NMR simulations on protein-size spin systems.

3 Simulation setting and experimental context

Fully ^{13}C and ^{15}N -labelled protein human ubiquitin (PDB code 1D3Z, Figure 1) containing over a thousand magnetic nuclei in 76 amino acid residues was cho-




Figure 1: Human ubiquitin protein (PDB code 1D3Z): 76 amino acids, 563 magnetic nuclei in the extended backbone (H, N, C, CA, CB, HA, HB)

sen for testing purposes with two types of spin system truncation: *backbone* (H, N, C, CA, HA) and *extended backbone* (H, N, C, CA, CB, HA, HB). Both cases involve a weakly branched continuous chain of spin-spin couplings and are encountered in the simulation of a large class of protein backbone NMR experiments — HNCO [26], HNCOCA [3], HNCA [26], etc. The isotropic NMR Hamiltonian was assembled using chemical shift values from the BMRB database [50] and J-couplings from the literature data [8, 37, 41, 52, 53]. In the cases where an experimental value of a particular J-coupling was not available in the literature, it was estimated based on the known values for structurally similar substances [2, 21, 33] — for most NMR simulation purposes and certainly for the purpose of the demonstration of the performance of the tensor train algorithm the accuracy of such coupling estimates (about 20%) is sufficient. The raw data for the magnetic couplings used in this work is available in the example set supplied with the current public version of the *Spinach* library [22].

NMR experiments were performed at 25°C on a Varian Inova 600MHz (14.1 Tesla) spectrometer equipped with a z-gradient triple-resonance cryogenic probe using a 0.5mM sample of uniformly ^{13}C and ^{15}N -labelled human ubiquitin in 10% D_2O . The ^{15}N spectra were collected as 2D ^1H - ^{15}N HSQC [4] spectra incorporating gradient enhanced coherence selection [25] and water flip-back. The spectra were recorded with acquisition times of 150ms (t_1 , ^{15}N) and 500ms (t_2 , ^1H). During the ^{15}N evolution period $^1\text{J}_{\text{HN}}$ and $^{1,2}\text{J}_{\text{NC}}$ couplings were either allowed to evolve, or decoupled by insertion of a rectangular ^{15}N or a shaped 200 μs ^{13}C inversion pulse using the central lobe of a sinc function. During ^1H acquisition ^{15}N nuclei were either evolved or decoupled using 40ppm broadband WURST sequence [34].

The liquid-state NMR Hamiltonian of ^{13}C , ^{15}N -labelled ubiquitin is:

$$\begin{aligned} \hat{H}(t) = & \sum_k \omega_k \hat{S}_Z^{(k)} + 2\pi \sum_{l>m} J_{\text{strong}}^{(l,m)} \hat{S}^{(l)} \cdot \hat{S}^{(m)} + 2\pi \sum_{p>q} J_{\text{weak}}^{(p,q)} \hat{S}_Z^{(p)} \hat{S}_Z^{(q)} \\ & + \omega_X(t) \sum_r \hat{S}_X^{(r)} + \omega_Y(t) \sum_s \hat{S}_Y^{(s)} \end{aligned} \quad (8)$$

where canonical NMR spectroscopy notation is used [17], k index runs over all nuclei, l and m indices run over pairs of nuclei that belong to the same isotope, p and q run over pairs of nuclei that belong to different isotopes, r and s runs over the nuclei influenced by radiofrequency control pulses, $\omega_X(t)$ and $\omega_Y(t)$ are time profiles of those pulses, ω_k are offset frequencies arising from the chemical shielding of the corresponding nuclei [43], $J_{\text{strong}}^{(l,m)}$ are “strong” NMR J-couplings [44], $J_{\text{weak}}^{(p,q)}$ are “weak” NMR J-couplings [17], and the spin operators $\hat{S}_X^{(n)}$, $\hat{S}_Y^{(n)}$, $\hat{S}_Z^{(n)}$ are defined by Eq. (2). In the case of extended ubiquitin backbone, Hamiltonian in Eq. (8) contains 563 shielding terms, 1840 coupling terms, and 1126 radiofrequency terms. All calculations reported below were performed by extending the functionality of *Spinach* library [22] to the tensor train formalism and interfacing it to *TT-Toolbox* [39] where appropriate.

Due to the abundance of complicated multi-pulse NMR experiments with time-dependent Hamiltonians [17], magnetic resonance simulations are generally carried out in the time domain. They always require long-term evolution trajectories with accurate phases (at least 100ms, much longer than the reciprocal Hamiltonian norm) for the density operator $\hat{\rho}(t)$ under the Liouville–von Neumann equation:

$$\begin{aligned} \frac{d}{dt} \hat{\rho}(t) &= -i [\hat{H}(t), \hat{\rho}(t)] + \hat{R} (\hat{\rho}(t) - \hat{\rho}_{\text{eq}}), \\ O(t) &= \langle \hat{O} | \hat{\rho}(t) \rangle = \text{Tr} [\hat{O}^\dagger \hat{\rho}(t)], \\ \hat{\rho}_{\text{eq}} &= \frac{\exp(-\hat{H}/k_B T)}{\text{Tr} \exp(-\hat{H}/k_B T)}. \end{aligned} \quad (9)$$

where \hat{R} is the relaxation superoperator ($T_{1,2}$ model with literature values for relaxation times [9] was used in the present work), $\hat{\rho}_{\text{eq}}$ is the thermal equilibrium state, and \hat{O} is the observable operator, usually a sum of \hat{S}_X or \hat{S}_Y operators on the spins of interest. In very simple cases where the Hamiltonian is not time dependent, the general solution to Eq. (9) can be written as:

$$O(t) = \langle \hat{O} | \exp[-i(\hat{H} + i\hat{R})t] | \hat{\rho}_0 \rangle, \quad (10)$$

where \hat{H} is the Hamiltonian commutation superoperator.

Direct time domain evaluation of this equation in tensor train format, either using explicit operator exponentiation or Krylov type propagation techniques that avoid it, does not appear to be possible — in all cases described by Eq. (8) the ranks in the tensor train expansion quickly grow beyond the capacity of modern computers. Increasing the singular value cut-off threshold at the representation compression stage leads to catastrophic loss of accuracy. Fortunately, there are simple

cases (most notably pulse-acquire 1D NMR spectroscopy) where amplitudes at only a few specific frequencies are actually required for the Fourier transform of Eq. (10), meaning that the problem can be reformulated as:

$$O(\omega) = -i \left\langle \hat{O} \left| \left(\hat{H} + i\hat{R} + \omega \mathbf{1} \right)^{-1} \right| \hat{\rho}_0 \right\rangle. \quad (11)$$

That is, to compute an observable in point ω in the frequency domain, we need to solve a linear system $\left(\hat{H} + i\hat{R} + \omega \mathbf{1} \right) |\hat{x}\rangle = |\hat{\rho}_0\rangle$. The formulation in Eq. (11) sacrifices a great deal of generality of Eq. (9) (simulation of arbitrary NMR pulse sequences is no longer possible), but it does serve as a stepping stone and enables the demonstration calculation presented below.

4 Tensor train algorithm for the summation and solution of linear systems

The DMRG algorithm has been originally proposed [55, 56] to find the “ground state” of a Hermitian matrix A by the minimization of the Rayleigh quotient $Q(x) = \langle x | A | x \rangle / \langle x | x \rangle$. The dynamical DMRG algorithm [23] was then developed to find solution of a linear system $A|x\rangle = |b\rangle$ with a Hermitian positive definite matrix A by the minimization of the “energy function” $J(x) = \langle x | A | x \rangle - 2\Re\langle x | b \rangle$. Besides the change of a target function, the algorithms are similar.

In DMRG algorithms the vector of interest x is sought in a form of a tensor train / MPS in Eq. (4). The minimization over all tensor train cores $x^{(n)}$ simultaneously is a highly nonlinear and complicated task. To make the procedure feasible, it is replaced by a sequence of optimizations carried over one core at a time:

$$x_\star^{(n)} = \arg \min_{x^{(n)}} J(\tau(x^{(1)}, \dots, x^{(n)}, \dots, x^{(N)})). \quad (12)$$

The MPS/TT format is linear in all sites $x^{(n)}$, which can be written as $x = X_{\neq n} x^{(n)}$, where the frame matrix $X_{\neq n}$ maps the parameters of the TT-core $x^{(n)}$ to the vector x . The linearity allows to rewrite Eq. (12) as $x_\star^{(n)} = \arg \min_{x^{(n)}} J_n(x^{(n)}) = A_n^{-1} b_n$, where J_n is the energy function for the *local problem* $A_n x^{(n)} = b_n$ with $A_n = X_{\neq n}^* A X_{\neq n}$ and $b_n = X_{\neq n}^* b$. Using the non-uniqueness of the tensor train format, one can always construct the representation with the unitary frame matrix $X_{\neq n}$, that guarantees the stability of the local problem. After the solution $x_\star^{(n)}$ is computed, we substitute $x^{(n)} := x_\star^{(n)}$ in the tensor train, and continue for $n = 1, \dots, N$, and then back and forth the chain.

The convergence of the described above *one-site* DMRG depends on the initial guess, in particular the TT ranks, because they remain the same during the sequence of updates defined by Eq. (12). This is a severe restriction, and additional methods are used to *adapt* the TT ranks during the computations. One way to do it is to replace the optimization over the single cores by the optimization over the pairs of neighboring cores, and then adapt the TT rank between them. Another possibility is to extend the search space by auxiliary directions. The first method of the latter type is the *corrected one-site* DMRG algorithm [57] which targets in addition to x a surrogate of the next Krylov vector Ax .

For the solution of linear systems, the *alternating minimum energy* (AMEn) algorithm was recently proposed [12, 13], which also uses an additional direction for the purpose of adaptivity. The local optimization step in AMEn is carried over one site only. To increase the TT ranks and improve the convergence, the TT blocks are expanded by auxiliary information, $x^{(n)} := \begin{bmatrix} x_\star^{(n)} & z^{(n)} \end{bmatrix}$. The *enrichment* $z^{(n)}$ introduces new directions in the subspace spanned by $X_{\neq n+1}$. A good choice of the enrichment is the component $z^{(n)}$ of the TT representation (exact or approximate) $z = \tau(z^{(1)}, \dots, z^{(N)})$ for the residual $z = b - Ax$. The proposed AMEn algorithm is as fast as one-site methods but is rank adaptive as the two-site DMRG algorithm, and demonstrates a comparable convergence rates. For the solution of a linear system $Ax = b$ with a Hermitian positive definite matrix, it has a proven global bound on the geometrical convergence rate.

In this work we use the AMEn algorithm for two purposes. First, we apply it to a system with a trivial matrix $A = \mathbf{1}$, but a complicated right-hand side b being a sum of many elementary tensors like the one in Eq. (2). This allows us to compress a Hamiltonian returned by the *Spinach* package from the CP format given by Eq. (8) into the TT format Eq. (4). The Hamiltonian is stretched into a vector, and the target functional $J(x) = \|x - b\|^2$ is a Frobenius-norm distance between a given Hamiltonian b and Hamiltonian x sought in the tensor train format. The one-site optimization in Eq. (12) is effectively the solution of the over-determined linear system $X_{\neq n} x^{(n)} = b$ by least squares. For the unitary frame matrix we have $x_\star^{(n)} = X_{\neq n}^* b$, and therefore the local optimization step is solved by the contraction of the frame matrix with the given Hamiltonian b . The enrichment step uses the low-rank approximation of the residual $z = b - x$, which is acquired by one-site DMRG optimization.

Second, we compute the NMR spectrum of a spin system by solving linear systems in Eq. (11). Since the matrix $\hat{H} + i\hat{R} + \omega\mathbf{1}$ is not Hermitian, we symmetrize the problem, and for $\hat{R} = \mu\mathbf{1}$ obtain

$$\Re \hat{p}(\omega) = -\mu \left(\hat{H}^* \hat{H} + 2\omega \hat{H} + (\omega^2 + \mu^2) \mathbf{1} \right)^{-1} \hat{p}_0. \quad (13)$$

This equation is solved by the AMEn algorithm at all points ω in a selected spectral domain.

5 Results and discussion

As discussed above, a major problem in the application of tensor train methods to magnetic resonance simulation of large systems is the calculation of lengthy sums involved in the construction of spin Hamiltonians and density matrices, and their compression into the TT format. Fig. 2 illustrates the performance of our proposed solution to this problem in the case of minimal (H, N, C, CA, HA) and extended (H, N, C, CA, HA, CB, HB) ubiquitin backbone spin systems. The storage requirements for the TT format in Eq. (4) depend on all TT ranks r_1, \dots, r_{N-1} , and are characterized by the *effective* TT rank r , defined by $Nr^2 = \sum_{n=1}^N r_{n-1} r_n$. It is immediately apparent from the left panels of Fig. 2 that the primary showstopper — rapid growth in the tensor train bond dimension — has been removed by the

spin system	number of terms	AMEn summation		RSS IK-1(4,2) basis	
		nnz, M	time, sec	nnz, M	time, sec.
bb	5490	3.0	423	6.8	9080
ext bb	10408	6.3	1765	23	67500

Table 1: The summation time and storage size for the isotropic part of Hamiltonian given by Eq. (8) for minimal (H, N, C, CA, HA) and extended (H, N, C, CA, HA, CB, HB) ubiquitin backbone spin systems (denoted as ‘bb’ and ‘ext bb’). The large number of terms are summed up by the AMEn algorithm to the TT format Eq. (4), and the reduced state space method [22] to the sparse matrix.

AMEn method: the effective TT ranks stay below 50 for the extended backbone and below 40 for the minimal backbone, well within the capability of modern desktop workstations. Since r^2 is less than the number of terms in the CP representation, the TT format with Nr^2 operators provides more compact storage than the CP format.

An alternative approach to AMEn is the *binary* summation, which sequentially adds and compresses the components pairwise. The TT ranks after the binary summation, however, rise up to several hundred and thereby make the solution of linear systems in Eq. (13) exceedingly difficult. It is clear from the right panels of Fig. 2 that the CPU time requirements of the AMEn summation compared to the binary summation are essentially the same, making the AMEn procedure clearly superior for all practical purposes. We also note that the summation of the Hamiltonian given by Eq. (8) to the TT format is more efficient than the sparse representation using the reduced state space (RSS) technique [22], see Table 1.

The resulting representation of the ubiquitin backbone spin Hamiltonian matrix is, up to the rounding error of the complex double precision arithmetic, exact. In magnetic resonance spectroscopy this is an unprecedented development — ubiquitin NMR simulation is currently just about feasible [16], with significant approximations and colossal computational resources. Tensor train representation is therefore a large step forward, even though Eq. (11) is not applicable to arbitrary NMR experiments.

The accuracy of the ^{15}N pulse-acquire NMR spectra, computed using Eq. (13) with the AMEn algorithm, is below 10^{-6} across the frequency interval, as demonstrated in Fig. 3. The deviations of the TT simulation from the RSS one result from the basis set effects in the restricted state space method, and the rank truncation threshold in the AMEn method.

Due to the intrinsically low sensitivity of liquid state ^{15}N protein NMR spectroscopy, it is not possible to record the experimental equivalent of Fig. 3 directly with a sufficient signal-to-noise ratio; we have therefore taken a somewhat longer route to the experimental validation of the tensor train simulation — Fig. 4 shows experimental proton-detected ^{15}N – ^1H HSQC spectra of ubiquitin, compared to the simulations obtained at the basis set limit of the restricted state space formalism [24]. Perfect agreement is apparent in both cases. This provides an experimental evidence to the accuracy of the restricted state space approach. The tensor train results in Fig. 3 can now be justified by comparison to the RSS approximation. It is clear from Fig. 3 that the TT formalism performs as intended.

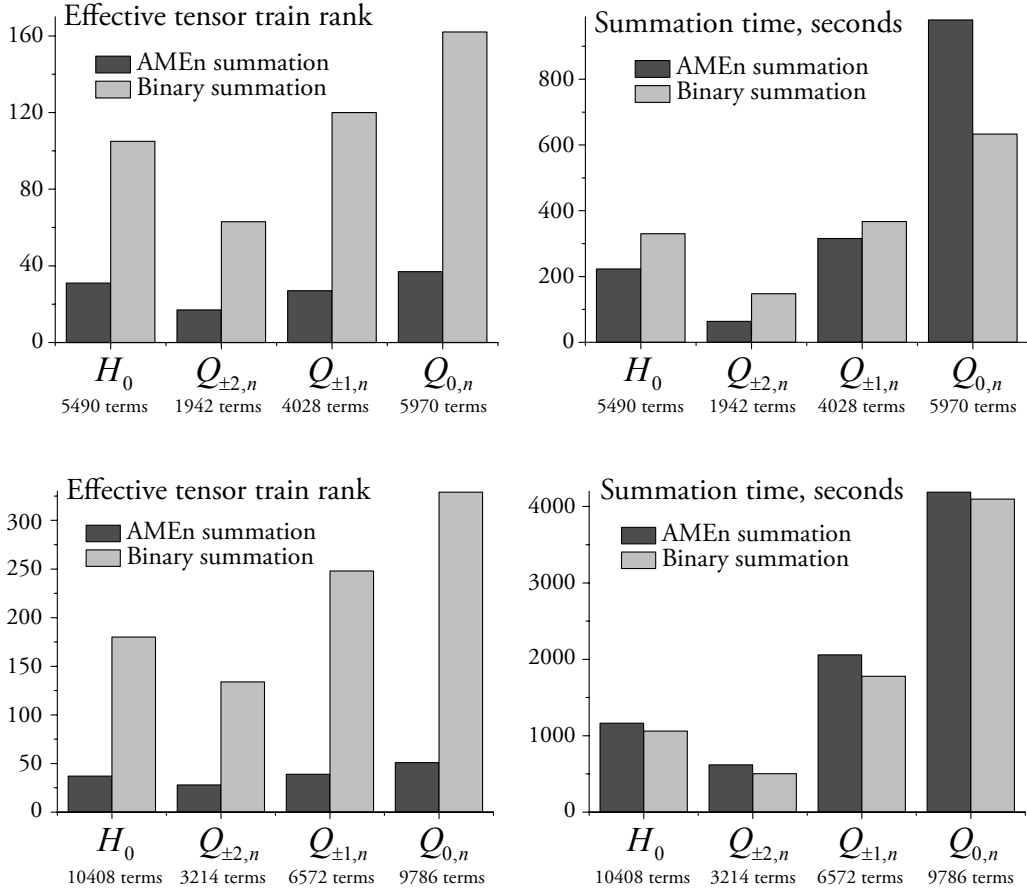


Figure 2: Performance comparison for binary and AMEn tensor train summation during the construction of the NMR spin Hamiltonian. **Top** — human ubiquitin backbone (H, N, C, CA, HA), **bottom** — human ubiquitin extended backbone (H, N, C, CA, HA, CB, HB). Here H_0 refers to the isotropic part of the Hamiltonian and $Q_{k,n}$ to the irreducible spherical components of the anisotropic part.

The successful 1D NMR simulation notwithstanding, very significant obstacles remain to the practical application of tensor train formalism to NMR spectroscopy. The following requirements should be addressed in future work to fully uncover the potential of the DMRG/MPS/TT formalism for spin dynamics simulation in NMR:

1. The DMRG requirement for the spin system to be a chain or a tree should be lifted. Biological magnetic resonance spin systems are irregular polycyclic interaction networks with multiple interlocking loops in the coupling graph, particularly in solid state NMR, where inter-nuclear dipolar couplings form a very dense network. A generalization of tensor product algorithms to complex tensor networks that fully mimics the molecule structure is required.
2. Rank explosion problem for time-domain simulations should be solved. It is clear from the success of the restricted state space approximation [35, 22] that the order of spin correlation in many evolving magnetic resonance spin systems either is or may safely be assumed to be quite low. This evidences data sparsity and indicates that some kind of low-rank decomposition is possible.




Figure 3: Amino group region of pulse-acquire ^{15}N NMR spectrum of human ubiquitin computed at the basis set limit of the restricted state space approximation [22] (upper trace) and using Eq. (11) as described in the main text. The relative deviation between the two spectra is below 10^{-6} across the frequency interval.

3. Our experience indicates that DMRG type objects are very far from being drop-in replacements for their matrix counterparts in standard simulation algorithms and software — it does actually appear that nearly everything in the very considerable body of magnetic resonance simulation methods needs to be adapted to the realities of DMRG. Current implementation of tensor product methods still requires a number of tuning parameters (e.g. approximation accuracies, TT ranks of the enrichment). The broad adoption of tensor product algorithms require the operations to be handled transparently and seamlessly by the existing simulation software packages, in the same way as sparse matrices currently are.
4. Transparent and clear tensor train approximation accuracy criteria, rank control and *a priori* error bounds should be developed — one of the most irksome factors we had encountered in our practical work with DMRG representations was the unpredictable accumulation of numerical rounding errors. This problem is particularly acute for the state vector phase in time domain simulations: magnetic resonance experiments rely critically on the phase being correctly predicted.

All of that having been said, we are still optimistic about the future of DMRG/MPS/TT methods, having also found them useful in Fokker–Planck type formalisms related to NMR and EPR spectroscopy [15]. Their primary strength is the lack of heuristic assumptions and the controllable nature of the representation accuracy. An experimental implementation of tensor train magnetic resonance simulation paths, via an interface to the *TT-Toolbox* [39], is available in version 1.3.1980 of our *Spinach* library [22].

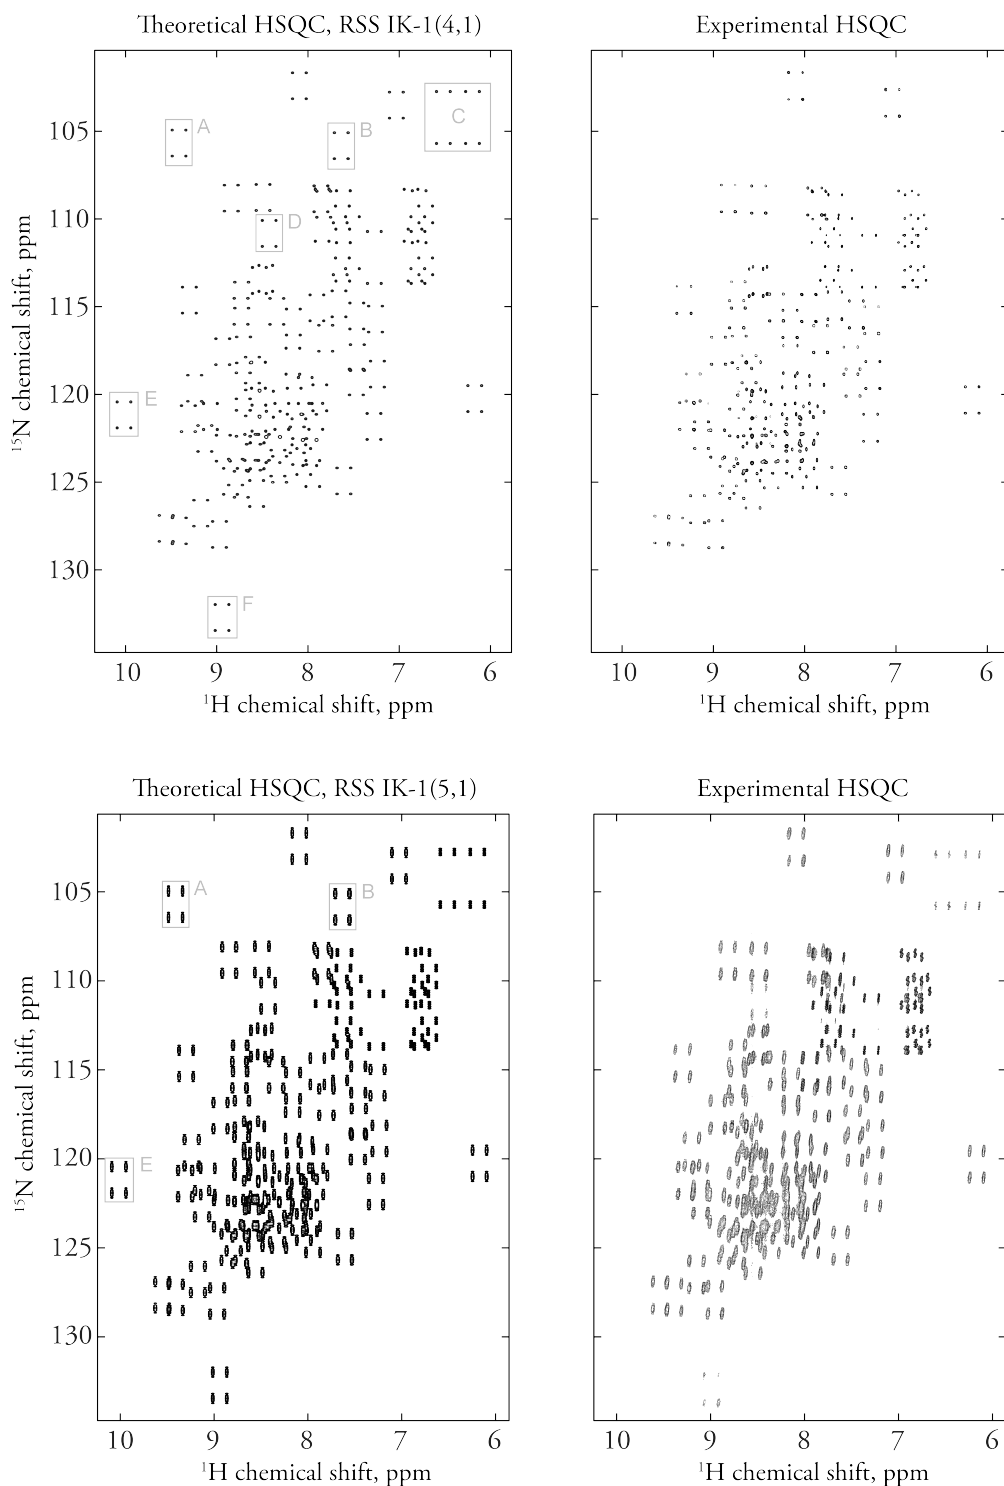


Figure 4: Theoretical (left) and experimental (right) ^1H - ^{15}N HSQC spectra of ^{15}N , ^{13}C -labelled human ubiquitin. **Top** — proton decoupling switched off in the indirect dimension and nitrogen decoupling switched off in the direct dimension to demonstrate the accurate quantum mechanical treatment of spin-spin coupling by the simulation. **Bottom** — additionally, carbon decoupling is also switched off in the indirect dimension. Signal groups marked A-F in the theoretical spectrum are not visible in the experimental data due to partial deuteration and slow conformational exchange of the corresponding amino acid residues.

6 Conclusions

Even with their well-documented limitations (the requirement for the spin system to be close to a chain, difficulty with long-range time-domain simulations, code implementation challenges, etc.), the ability of DMRG type formalisms to simulate simple liquid state NMR spectra of large spin systems essentially without approximations is impressive. They cannot yet match the highly optimized dedicated methods developed by the magnetic resonance community [22], but if some of the limitations are lifted by the subsequent research, tensor network methods would have the potential to become a very useful formalism in NMR research. Benefits to EPR spectroscopy, with its star-shaped spin interactions graphs, are likely to be harder to achieve, but may still be obtained by exploiting the direct product structure of combined spin and spatial dynamics appearing in Fokker-Planck type problems [15].

Having solved in this paper the last purely technical problem on the way to the broad adoption of tensor train formalism in magnetic resonance spectroscopy, we are quite optimistic about its potential. We look forward to generalizing the AMEn method to arbitrary tensor networks, that closely match the geometry of a molecule. We also plan to develop reliable methods for indefinite matrices, as well as time evolution schemes maintaining conservation laws in tensor formats.

Acknowledgements

We are grateful to Garnet Chan for patiently explaining the DMRG formalism to IK during his visit to Cornell University and to Zenawi Welderufael for finding some of the less obvious ubiquitin J-couplings in the literature. We appreciate the use of the Iridis High Performance Computing Facility, and the associated support services at the University of Southampton, that proved essential to carry out the extensive numerical experiments reported in this paper. JMW would like to acknowledge the Wellcome Trust for support of the Southampton NMR centre and we would like to thank the Geoff Kelly at NIMR (Mill Hill) for lending us the ubiquitin sample. The project is supported by EPSRC (EP/H003789/2, EP/J013080/1).

References

- [1] M. BAK, J. T. RASMUSSEN, AND N. C. NIELSEN, *SIMPSON: A general simulation program for solid-state NMR spectroscopy*, J Magn. Reson., 147 (2000), pp. 296–330.
- [2] M. BARFIELD, R. J. SPEAR, AND S. STERNHELL, *Allylic interproton spin–spin coupling*, Chem. Rev., 76 (1976), pp. 593–624.
- [3] A. BAX AND M. IKURA, *An efficient 3D NMR technique for correlating the proton and ^{15}N backbone amide resonances with the α -carbon of the preceding residue in uniformly $^{15}\text{N}/^{13}\text{C}$ enriched proteins*, Journal of biomolecular NMR, 1 (1991), pp. 99–104.
- [4] G. BODENHAUSEN AND D. J. RUBEN, *Natural abundance nitrogen-15 NMR by enhanced heteronuclear spectroscopy*, Chemical Physics Letters, 69 (1980), pp. 185–189.
- [5] R. BRÜSCHWEILER AND R. R. ERNST, *A cog-wheel model for nuclear-spin propagation in solids*, J. Magn. Reson., 124 (1997), pp. 122–126.
- [6] R. BURSILL, T. XIANG, AND G. GEHRING, *The density matrix renormalization group for a quantum spin chain at non-zero temperature*, Journal of Physics: Condensed Matter, 8 (1996), p. L583.
- [7] M. C. BUTLER, J. N. DUMEZ, AND L. EMSLEY, *Dynamics of large nuclear-spin systems from low-order correlations in Liouville space*, Chem. Phys. Lett., 477 (2009), pp. 377–381.
- [8] D. A. CASE, C. SCHEURER, AND R. BRÜSCHWEILER, *Static and dynamic effects on vicinal scalar J couplings in proteins and peptides: a MD/DFT analysis*, J. Am. Chem. Soc., 122 (2000), pp. 10390–10397.
- [9] J. CAVANAGH, W. J. FAIRBROTHER, A. G. PALMER III, AND N. J. SKELTON, *Protein NMR spectroscopy: principles and practice*, Elsevier Academic Press, 1995.
- [10] S. V. DOLGOV, B. N. KHOROMSKIJ, AND D. V. SAVOSTYANOV, *Superfast Fourier transform using QTT approximation*, J. Fourier Anal. Appl., 18 (2012), pp. 915–953.
- [11] S. V. DOLGOV AND I. V. OSELEDETS, *Solution of linear systems and matrix inversion in the TT-format*, SIAM J. Sci. Comput., 34 (2012), pp. A2718–A2739.
- [12] S. V. DOLGOV AND D. V. SAVOSTYANOV, *Alternating minimal energy methods for linear systems in higher dimensions. Part I: SPD systems*, arXiv preprint 1301.6068, 2013.
- [13] —, *Alternating minimal energy methods for linear systems in higher dimensions. Part II: Faster algorithm and application to nonsymmetric systems*, arXiv preprint 1304.1222, 2013.
- [14] R. S. DUMONT, S. JAIN, AND A. BAIN, *Simulation of many-spin system dynamics via sparse matrix methodology*, J Chem. Phys., 106 (1997), pp. 5928–5936.

- [15] L. E. EDWARDS, D. V. SAVOSTYANOV, A. A. NEVZOROV, M. CONCISTRÈ, G. PILEO, AND I. KUPROV, *Grid-free power averages: On the application of Fokker-Planck equation to solid state NMR*, J. Magn. Reson., 235 (2013), pp. 121–129.
- [16] L. J. EDWARDS, D. SAVOSTYANOV, Z. WELDERUFAEL, D. LEE, AND I. KUPROV, *A quantum mechanical NMR simulation algorithm for protein-scale spin systems*, J. Am. Chem. Soc.
- [17] R. R. ERNST, G. BODENHAUSEN, AND A. WOKAUN, *Principles of nuclear magnetic resonance in one and two dimensions*, vol. 14, Clarendon Press Oxford, 1987.
- [18] M. FANNES, B. NACHTERGAELE, AND R. WERNER, *Finitely correlated states on quantum spin chains*, Communications in Mathematical Physics, 144 (1992), pp. 443–490.
- [19] S. A. GOREINOV, I. V. OSELEDETS, AND D. V. SAVOSTYANOV, *Wedderburn rank reduction and Krylov subspace method for tensor approximation. Part 1: Tucker case*, SIAM J. Sci. Comput., 34 (2012), pp. A1–A27.
- [20] W. HACKBUSCH, *Tensor spaces and numerical tensor calculus*, Springer–Verlag, Berlin, 2012.
- [21] P. E. HANSEN, *Carbon–hydrogen spin–spin coupling constants*, Progress in Nuclear Magnetic Resonance Spectroscopy, 14 (1981), pp. 175–295.
- [22] H. J. HOGBEN, M. KRZYSTYNIAK, G. T. P. CHARNOCK, P. J. HORE, AND I. KUPROV, *Spinach — A software library for simulation of spin dynamics in large spin systems*, J Magn. Reson., 208 (2011), pp. 179–194.
- [23] E. JECKELMANN, *Dynamical density–matrix renormalization–group method*, Phys Rev B, 66 (2002), p. 045114.
- [24] A. KARABANOV, I. KUPROV, G. T. P. CHARNOCK, A. VAN DER DRIFT, L. J. EDWARDS, AND W. KÖCKENBERGER, *On the accuracy of the state space restriction approximation for spin dynamics simulations*, J Chem. Phys., 135 (2011), p. 084106.
- [25] L. KAY, P. KEIFER, AND T. SAARINEN, *Pure absorption gradient enhanced heteronuclear single quantum correlation spectroscopy with improved sensitivity*, J. Am. Chem. Soc., 114 (1992), pp. 10663–10665.
- [26] L. E. KAY, M. IKURA, R. TSCHUDIN, AND A. BAX, *Three-dimensional triple-resonance NMR spectroscopy of isotopically enriched proteins*, J. Magn. Reson. (1969), 89 (1990), pp. 496–514.
- [27] V. KAZEEV, B. KHOROMSKIJ, AND E. TYRTYSHNIKOV, *Multilevel Toeplitz matrices generated by tensor-structured vectors and convolution with logarithmic complexity*, SIAM J. Sci. Comp., 35 (2013), pp. A1511–A1536.
- [28] V. A. KAZEEV AND B. N. KHOROMSKIJ, *Low-rank explicit QTT representation of the Laplace operator and its inverse*, SIAM J. Matrix Anal. Appl., 33 (2012), pp. 742–758.

- [29] B. N. KHOROMSKIY, *Tensor-structured numerical methods in scientific computing: Survey on recent advances*, Chemometr. Intell. Lab. Syst., 110 (2012), pp. 1–19.
- [30] B. N. KHOROMSKIY AND V. KHOROMSKAIA, *Multigrid accelerated tensor approximation of function related multidimensional arrays*, SIAM J. Sci. Comput., 31 (2009), pp. 3002–3026.
- [31] A. KLÜMPER, A. SCHADSCHNEIDER, AND J. ZITTARTZ, *Matrix product ground states for one-dimensional spin-1 quantum antiferromagnets*, Europhys. Lett., 24 (1993), pp. 293–297.
- [32] T. G. KOLDA AND B. W. BADER, *Tensor decompositions and applications*, SIAM Review, 51 (2009), pp. 455–500.
- [33] L. B. KRIVDIN AND E. W. DELLA, *Spin–spin coupling constants between carbons separated by more than one bond*, Progress in Nuclear Magnetic Resonance Spectroscopy, 23 (1991), pp. 301–610.
- [34] Ē. KUPCE AND R. FREEMAN, *Stretched adiabatic pulses for broadband spin inversion*, Journal of Magnetic Resonance, Series A, 117 (1995), pp. 246–256.
- [35] I. KUPROV, N. WAGNER-RUNDELL, AND P. J. HORE, *Polynomially scaling spin dynamics simulation algorithm based on adaptive state-space restriction*, J Magn. Reson., 189 (2007), pp. 241–250.
- [36] S. LIANG AND H. PANG, *Approximate diagonalization using the density matrix renormalization-group method: A two-dimensional-systems perspective*, Phys. Rev. B, 49 (1994), pp. 9214–9217.
- [37] F. LÖHR, J. M. SCHMIDT, AND H. RÜTERJANS, *Simultaneous measurement of $^3J_{\text{HN,H}\alpha}$ and $^3J_{\text{H}\alpha,\text{H}\beta}$ coupling constants in ^{13}C , ^{15}N -labeled proteins*, J. Am. Chem. Soc., 121 (1999), pp. 11821–11826.
- [38] I. P. MCCULLOCH, *Density matrix renormalization group algorithm and the two-dimensional t-J model*, Philosophical Magazine Part B, 81 (2001), pp. 1603–1613.
- [39] I. OSELEDETS ET AL., *TT-Toolbox*.
- [40] I. V. OSELEDETS, *Tensor-train decomposition*, SIAM J. Sci. Comput., 33 (2011), pp. 2295–2317.
- [41] C. PÉREZ, F. LÖHR, H. RÜTERJANS, AND J. M. SCHMIDT, *Self-consistent Karplus parametrization of 3J couplings depending on the polypeptide side-chain torsion ξ_1* , J. Am. Chem. Soc., 123 (2001), pp. 7081–7093.
- [42] D. PEREZ-GARCIA, F. VERSTRAETE, M. M. WOLF, AND J. I. CIRAC, *Matrix product state representations*, Quantum Info. Comput., 7 (2007), pp. 401–430.
- [43] N. F. RAMSEY, *Magnetic shielding of nuclei in molecules*, Physical Review, 78 (1950), p. 699.
- [44] ———, *Electron coupled interactions between nuclear spins in molecules*, Physical Review, 91 (1953), p. 303.

- [45] D. V. SAVOSTYANOV, *Fast revealing of mode ranks of tensor in canonical form*, Numer. Math. Theor. Meth. Appl., 2 (2009), pp. 439–444.
- [46] U. SCHOLLWÖCK, *The density–matrix renormalization group*, Rev. Mod. Phys., 77 (2005), pp. 259–315.
- [47] ———, *The density-matrix renormalization group in the age of matrix product states*, Annals of Physics, 326 (2011), pp. 96–192.
- [48] N. SHIBATA, *Thermodynamics of the anisotropic heisenberg chain calculated by the density matrix renormalization group method*, J. Phys. Soc. Jpn., 66 (1997), pp. 2221–2223.
- [49] E. SToudenMIRE AND S. R. WHITE, *Studying two-dimensional systems with the density matrix renormalization group*, Annual Review of Condensed Matter Physics, 3 (2012), pp. 111–128.
- [50] E. L. ULRICH, H. AKUTSU, J. F. DORELEIJERS, Y. HARANO, Y. E. IOANNIDIS, J. LIN, M. LIVNY, S. MADING, D. MAZIUK, Z. MILLER, E. NAKATANI, C. F. SCHULTE, D. E. TOLMIE, R. K. WENGER, H. YAO, AND J. L. MARKLEY, *Biomagresbank*, Nucleic acids research, 36 (2008), pp. D402–D408.
- [51] M. VESHTORT AND R. G. GRIFFIN, *SPINEVOLUTION: A powerful tool for the simulation of solid and liquid state NMR experiments*, J Magn. Reson., 178 (2006), pp. 248–282.
- [52] B. VÖGELI, J. YING, A. GRISHAEV, AND A. BAX, *Limits on variations in protein backbone dynamics from precise measurements of scalar couplings*, J. Am. Chem. Soc., 129 (2007), pp. 9377–9385.
- [53] A. C. WANG AND A. BAX, *Determination of the backbone dihedral angles ϕ in human ubiquitin from reparametrized empirical karplus equations*, J. Am. Chem. Soc., 118 (1996), pp. 2483–2494.
- [54] X. WANG AND T. XIANG, *Transfer-matrix density-matrix renormalization-group theory for thermodynamics of one-dimensional quantum systems*, Phys. Rev. B, 56 (1997), pp. 5061–5064.
- [55] S. R. WHITE, *Density matrix formulation for quantum renormalization groups*, Phys. Rev. Lett., 69 (1992), pp. 2863–2866.
- [56] ———, *Density-matrix algorithms for quantum renormalization groups*, Phys. Rev. B, 48 (1993), pp. 10345–10356.
- [57] ———, *Density matrix renormalization group algorithms with a single center site*, Phys. Rev. B, 72 (2005), p. 180403.
- [58] T. XIANG, J. LOU, AND Z. SU, *Two-dimensional algorithm of the density-matrix renormalization group*, Phys. Rev. B, 64 (2001), p. 104414.